Don't listen, talk:
A comparative study of transmit and receive power for low power radios.

A Greenhalgh†and S Hailes†

†Dept. of Computer Science, University College London

Abstract

In this paper we investigate the energy consumption of a low power radio device in its three modes of operation: transmission, reception and idle. We identify these modes of operation from the traces of the energy consumption of the device by modifying the device's firmware to include and exclude certain events. We present energy consumption measurements two orders of magnitude lower than an 802.11 device and question whether the use of 802.11 figures in simulations of low power radio systems is realistic. Contrary to 802.11 models, we show $P_{rx} > P_{tx} > P_{idle}$.

1 Introduction

Building a large scale ad hoc network testbed to test and model ideas about new protocols, device designs and applications is too expensive and infeasible for all but the very largest research budgets. Simulation of ad hoc networks is the approach taken by many when researching ad hoc networks; however, simulations are only as accurate as the models and parameters used in them. This paper aims to improve the energy consumption model of the simulator by measuring the important parameters of an actual low power radio device. When simulating low power radio devices for use in ad hoc networks three key parameters are used to calculate the energy consumption of the device: the energy consumption when transmitting, P_{tx}, the energy consumption when receiving data P_{rx} and the energy consumption when the device is idle (i.e. when not transmitting or receiving data), P_{idle}.

1.1 Related work

Related work has been carried out by [1], [2], [3] and [4], but this work has focused on the 802.11b family of devices. None of these pieces of work has had access to the device firmware and none have been able to carry out an event analysis of the energy consumption traces as we do here. The techniques used here to calculate the energy consumption of different stages of the operation are similar to those used by Feeney et al. [1]; however we take these techniques a stage further by disabling specific sections of the firmware to aid identification.

2 Experimental setup

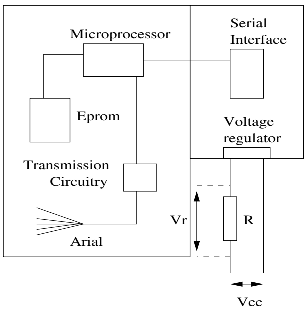

Figure 1: Experimental Configuration

Given the nature of the device, and low power radio devices in general, it is infeasible to measure the power consumption of specific components. It is, however, simple to measure the overall power consumption of the device. The overall power consumption of the device is , $P = I_{cc}V_{cc}$, where Vcc is the voltage across inputs to the circuit. [1] and I_{cc} is the current flowing through the device, shown in figure 1. I_{cc} can be measured by measuring the voltage, V_r, across a resistor, R, placed in line with the supply V_{cc}, $I_{cc} = V_r/R$. As the device performs different functions, computation, transmission, reception, etc. I_{cc} varies.

Using a digital storage oscilloscope to measure and store V_r, the power consumption of a specific event starting at time t_0 and finishing at time t_1 can be measured by applying equation 1 .

$$P_{t_0,t_1} = \frac{V_{cc}}{R} \sum_{t=t_0}^{t_1} V_{r_t} \quad (1)$$

Equation 1 assumes that both V_{cc} and R are constant with respect to time throughout the experiment. Figure 1 shows the experimental configuration.

[1]The positive voltage input is called the Common Collector (cc) and V_{cc} is measure between the common collector and ground (GND)

Three versions of the device firmware were developed; a version that sources packets (a source), a version that receives packets (a sink) and a version that receives packets and retransmits them (a relay). The source device can be configured to generate packets of 9, 39, 69 or 99 bytes continually at a rate of 5 packets per second, with its transmitter circuitry either enabled or disabled. The devices are used in two different scenarios: individually and in pairs. Each experiment was repeated 5 times and the results averaged. Where specific events need to be identified in each experiment, such the start of a transmission, these are carried out individually for each iteration of an experiment.

A custom development board developed by Philips Research Labs (Redhill) for low power radio applications was modified to remove all software layers except the Medium Access Control layer (MAC) and Physical layers to enable the power consumption of only the radio interface software and hardware to be measured. The development board has an 8-bit Mitsubishi 3807, a serial interface, a 512Kbit eprom and a RFM TR1001L radio module. The MAC collision avoidance scheme in use is CDMA/CA. The digital storage scope used here is a HP Model 5450C.

3 Single device experiments

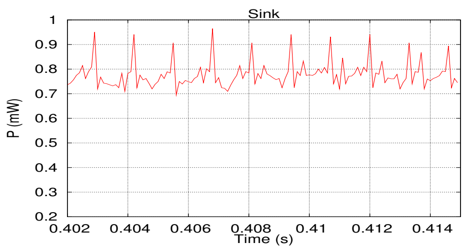

Figure 2: P (mW) consumption for a sink

Here we aim to find the quiescent energy consumption costs of the sink, relay and source devices in their different phases of operation (transmitting data, awaiting or receiving data and idle). A sink and relay device were run separately in the absence of any other device and P measured. Due to the limited space, we only show the energy consumption of a sink device, in figure 2. The exact causes of any individual peak in the energy consumption of the sink are not known, but we speculate that the large periodically repeating peaks, for example at 0.412s, are caused by they system oscillator.

The oscillator for these devices operates at 7.3728Mhz, which gives a peak-to-peak distance of $1.37 * 10^{-4}$ seconds which is approximately the peak-to-peak distance seen in figure 2. Table 1 shows the average energy consumption of a receiver and a relay over the full period of the experiment.

Device	Idle	Rx	Tx
Sink	NA	0.78	NA
Relay	NA	0.79	NA

Table 1: Power consumption of a sink and relay, P(mW)

3.1 Source

Transmitter	Enabled			Disabled		
Bytes	Idle	Rx	Tx	Idle	Rx	Tx
09	0.22	0.78	0.7	NA	0.78	NA
39	0.21	0.77	0.68	NA	0.79	NA
69	0.23	0.79	0.72	NA	0.78	NA
99	0.24	0.81	0.75	NA	0.79	NA

Table 2: P(mW) in the different phases

The source firmware allows for the transmitter circuitry to be disabled, isolating the costs due to the transmission. This can be seen in figures 3(a) (in which the transmitter is disabled) and 3(b) (in which the transmitter is enabled). By comparing figures 3(a) and 3(b), the transmission can be isolated: it occurs between 0.406s and 0.413s in 3(b). Before each transmission, the receiver circuitry is switched off and the device transitions to the idle state. Having isolated the transmission, the idle phase is from 0.403s to 0.406s, in figure 3(a). Table 2 shows the costs of the device in each of its different phases of operation; idle, receiving, and transmitting.

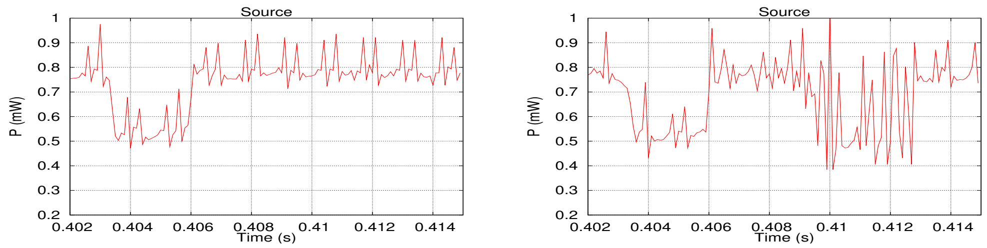

(a) transmitter disabled

(b) transmitter enabled

Figure 3: Energy consumption of a source, P(mW), 4 byte packet

Packet size (bytes)	Idle	Rx	Tx
9	0.22	0.77	0.69
39	0.21	0.77	0.68
69	0.22	0.77	0.71
99	0.22	0.78	0.71

Table 3: Power consumption of a source when transmitting in the presence of a relay, P(mW)

Packet size (bytes)	Idle	Rx	Tx
9	0.23	0.79	0.71
39	0.22	0.79	0.71
69	0.22	0.78	0.74
99	0.23	0.78	0.74

Table 4: Power consumption of a relay when re-transmitting in the presence of a source, P(mW)

3.2 Summary of results

The cost of reception was found to be 0.79 mW (SD 0.01), transmission was found to be 0.73 mW (SD 0.03) and the cost of being idle was found to be 0.23 mW (SD 0.01). Inspection of the data implied that there might be a relationship between the length of the transmission and the energy consumption.

4 Two device experiments

The previous section reported isolated tests of the devices; in this section, we test the interaction of the devices. Specifically, we look at the interactions between a source and a relay, and a source with a sink or a relay. Because sinks do not transmit packets they cause no interactions with sources, hence we do not need to measure the effect of a sink on a source. This is not necessarily the case in general for low power radio devices.

4.1 Source and Relay

In this experiment, a source sends data to a relay device that retransmits whatever it hears. The experiment is carried out four times using different packets sizes (9, 39, 69, and 99 bytes). Table 3 shows the energy consumption of the source device and table 4 shows the energy consumption of the relay device.

Comparing the cost of reception for the two device experiments with the cost of reception for the single device experiment, table 1, shows the energy consumption of the source devices and relay devices are approximately the same whether there is a transmitting source or not. The energy consumption of a relay device is, on average, higher than that of a source device in both scenarios.

Transmission costs are approximately 90% of the cost of reception. This is interesting because this is result is different to the assumptions of many schemes where transmission costs are expected to be higher than reception costs. We briefly discuss the implications of this later in section 5.

4.2 Source and Sink

One of the challenges faced whilst conducting this set of experiments was in measuring the effect of a transmission upon a receiving device. It was found that no identifiable variation in V_r in either the amplitude or frequency domains is observable to be used as a trigger for the oscilloscope. The oscilloscope's built-in Fast Fourier Transform routine failed to identify a frequency change. A possible reason for this is that the cost of decoding a "signal" is constant whether a signal is present or not.

4.3 Summary of results

The cost of reception was found to be 0.78 mW (SD 0.01), transmission was found to be 0.71 mW (SD 0.02) and the cost of being idle was found to be 0.22 mW (SD 0.01). Comparison with the single device experiments does not indicate any influence from the presence of additional devices on the energy consumption of the devices. This is not the case in general, but the firmware used here has no flow control in the MAC layer, or any higher layer. Introducing an 802.11-style CTS/RTS flow control scheme would introduce inter-device dependencies leading to costs related to the presence of other devices.

4.4 Limitations and sources of error

No thought was given to the effect of collisions between packets upon the energy costs; this is something that was outside the scope of this experiment and is complex to measure because it was not possible to determine the start of packet reception from the energy traces.

The known sources of error include the value of the resistor, which was rated with a 2% tolerance, and the identification of the key points within the voltage V_r trace. If the points are identified to an accuracy of ± 2 points on either side, this leads to a 0.002% inaccuracy for a data size of 8192 points. The unknown sources of error include the voltage drift of V_{cc} from recorded value at the start of the experiment, any device variation, and any errors associated with the oscilloscope. The effect of the unknown sources of error is thought to be minimal and unlikely to affect the results, whilst the effect of the known sources of error is thought to have a minimal affect on the results, given a maximum observed SD of 0.02.

5 Conclusions

In this paper, we have investigated a real low power radio system and gathered real energy consumption parameters for use in large-scale simulation. There are two key findings. Firstly, the use of 802.11 based parameters as a model for low power networks is unrealistic (table 5) showing that the energy consumption of the low power radio device measured here is three orders of magnitude lower than the 802.11 devices measured in the Feeney study [1].

Experiment	Feeney [1]		UCL
Mode	2Mbps	11 Mbps	19.2Kbps
Idle	843.72 mW	739.44 mW	0.22mW
Receive	966.96 mW	900.60 mW	0.78mW
Transmit	1327.20 mW	1346.16 mW	0.71mW

Table 5: Comparison of 802.11 and a low power radio device energy consumption figures

Secondly, the cost of reception is higher than the cost of transmission. One possible explanation for this is that the cost of the signal processing to decode the radio signal is more complex than encoding of the signal for transmission. However, the important observation here is that the typical assumption that $P_{tx} > P_{rx} > P_{idle}$ does *not* hold for this low power radio system; rather $P_{rx} > P_{tx} > P_{idle}$. This result poses an interesting question: is a MAC protocol that keeps the radio listening to the channel continually, such as CSMA/CA, the best MAC protocol for low power radio systems? Would a slotted time-based protocol be better a better fit, perhaps TDMA?

Acknowledgments.

The authors wish to acknowledge the considerable support of the Low Power Radio group at Philips Research Labs, Redhill, through their provision of the hardware and software for the experiment. Adam Greenhalgh was funded through an EPSRC Industrial CASE award with Philips Research Labs, Redhill.

Contact Address

The authors can be contacted by email; (a.greenhalgh,s.hailes)@cs.ucl.ac.uk.

References

[1] Laura Marie Feeney, Martin Nilsson *Investigating the Energy Consumption of a Wireless Network Interface in an ad hoc Networking Environment*Infocom 2001

[2] Mark Stemm and Randy H. Katz.*Measuring and reducing energy consumption of network interfaces in hand-held devices*

[3] Robin Kravets, Karsten Schwan, Ken Calvert.*Power-Aware Communication for Mobile Computers*

[4] Robin Kravets, P. Kirshnan.*Power Management Techniques for mobile communication*